



Relationship of OSG and NRP

**Frank Würthwein
Director
San Diego Supercomputer Center**

March 16th 2022



Shared Vision & Aspiration





Democratizing Access



The Minds We Need

- **Connect every community college, every minority serving institution, and every college and university, including all urban, rural, and tribal institutions** to a world-class and secure R&E infrastructure, with particular attention to institutions that have been chronically underserved;
- **Engage and empower every student and researcher** everywhere with the opportunity to join collaborative environments of the future, because we cannot know where the next Edison, Carver, Curie, McClintock, Einstein, or Katherine Johnson will come from; and

<https://mindsweneed.org>



Long Term Vision



- Create an Open National Cyberinfrastructure that allows the federation of CI at all ~4,000 accredited, degree granting higher education institutions, non-profit research institutions, and national laboratories.

- Open Science

- Open Data

- Open Source

- Open Infrastructure

Open Compute

Open Storage & CDN

Open devices/instruments/IoT, ...?

Openness for an Open Society



3 Ways to build Open Infrastructure Federation(s)



- At the container orchestration layer, e.g. K8S as implemented in Pacific Research Platform.
- Federating independent container orchestration frameworks, e.g. via SLATE or via Admiralty on K8S.
- Federating storage and compute clusters at the storage and batch system layer., e.g. via a traditional OSG setup with a CE.

**More (cybersecurity) control
implies more effort to join & operate.**

Some Institutions will never be able to join because of a mismatch between the effort they have and the control they desire.



Complementarity in Implementation

OSG focused on campus cluster integration.
Pacific Research Platform focused on individual node integration instead of clusters.

PRP is a single K8S cluster across the globe.

(Though, PRP has done federation across K8S clusters using admiralty with both Google and Expanse, and may do more of that in the future.)

A bit of history

- **Pacific Research Platform** funded by NSF CC*DNI DIBBs from 10/1/2015 to 9/30/2022 (PI Larry Smarr)
 - Developed core technologies (NSF 1541349)
- **CHASE-CI** funded by NSF as part of CI-New from 10/1/2017 to 9/30/2021 (PI Larry Smarr)
 - Provided about half to GPU hardware (NSF 1730158)
- Multiple other awards & community investments provided additional hardware and effort.
- Success on west coast lead to global growth.

**Separate OAC and CISE investments
=> Substantial CS machine learning community**

(See other contributing awards on last slide)



Pacific Research Platform



- “Bring Your Own Resource” philosophy same as OSG
 - From “soup kitchen” to “potluck” supercomputing
- Pioneered integration at the K8S & IPMI layer
 - Integrating resources at the single node level
 - Provide updated hardware shopping list for institutions
 - 2021: 11 different options from 8-way GPU, to 4-way FPGA, to 648TB SAS, ...
 - Among the most popular are 8-way GPU and 648TB Ceph storage nodes
 - join at K8S layer (institutions responsible for their own base OS install)
 - Or allow central team to devOps base OS via IPMI from UCSD/UNL
- Global hardware integration across 30++ institutions
 - 11 MSIs (8 in CA) and 6 institutions in EPSCoR states
- User interfaces: Jupyter, kubectl, OSG, HTCondor
- Community dominated by CS and EE
 - 500++ GPUs & many PBs of disk space for machine learning



PRP meets OSG



- OSG operates two K8S clusters in fixed locations, Wisconsin Madison & U.Chicago
 - (see Lincoln's talk Wednesday afternoon)
 - All centrally deployable containerized services of OSG are operated in one of these two K8S clusters.
- **OSG Data Origins and Data Caches are deployed via containers on PRP integrated hardware worldwide**
 - Institution joins hardware into OSG namespace in PRP K8S cluster & PATH team operates services.
- The PRP K8S cluster appears in OSG as a compute resource: SDSC-PRP
 - OSG users can thus use GPUs on PRP.



PRP as OSG Resource Provider



	total
IceCube	899 K
COVID19_FoldingAtHome	173 K
nova	16 K
PixleyLab	9 K
LIGO	7 K
TG-CHE200122	4 K
COVID19_JHU_Howard	2 K
ligo	544
JHU_Howard	134
FSU_Kolberg	8
TG-DDM160003	8
NeuroscienceGateway	1.3
COVID19_UCSD_Hsiao	1.2
OSG-Staff	0.9
USC_Deelman	0.9
fermilab	0.14
IIT_Cheng	0.12
CampusWorkshop_Feb2021	0.11
osg.NeuroscienceGateway	0.11
GATech_PACE	0.057

**More than 1 Million GPU hours
on PRP used via OSG integration
within the last 2 years**



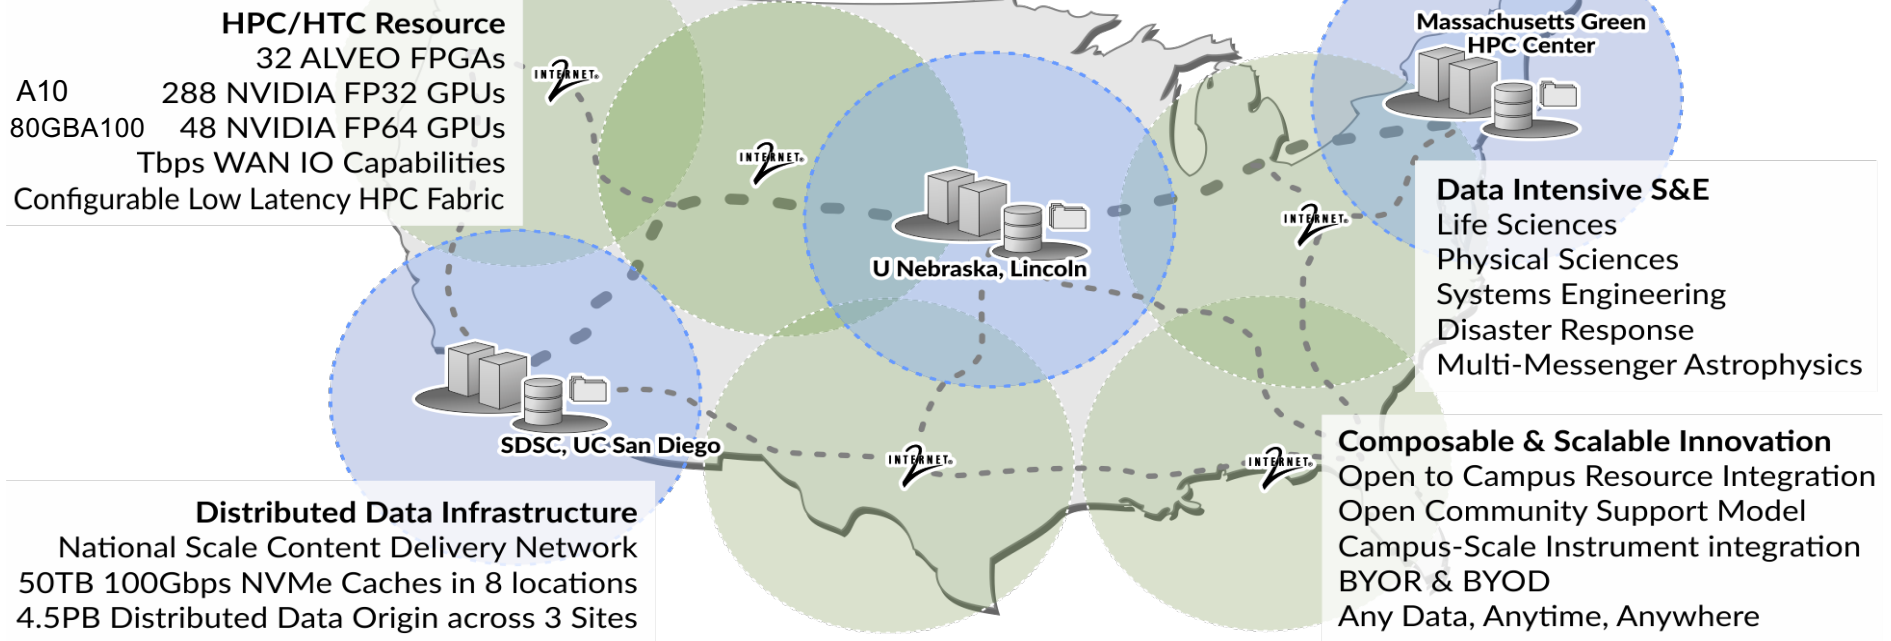


New in 2022

Systematically cover the entire continental united states

NATIONAL RESEARCH PLATFORM

Designed for Growth & Inclusion



Nationally Distributed Category-II System (NSF)

3 year prototyping phase followed by 2 year allocation phase.



Data Infrastructure Model of NRP



- **Support regional Ceph storage systems across the USA.**
 - Campuses can join individual storage hosts to the Ceph system in their region.
 - All regional storage systems are Origins in OSG Data Federation (OSDF)
 - Deploy replication system such that researchers can decide what part of their namespace should be in which regional storage.
- **Deploy caches in Internet2 backbone such that no campus nationwide is more than 500 miles from a cache.**
- Between all the collaborating projects, we ought to have >10PB of data origin storage in OSDF to support open science by the end of 2022.

NRP is a much more experimental data infrastructure than OSDF.

Filesystem mounts of regional storage on compute in the region.

User controlled replication of partial namespaces across region.

Institutional BYOR into regional storage systems.

NRP caches are part of the OSDF.



Summary & Conclusion



- It takes a cyberinfrastructure ecosystem of collaborating projects to build and operate the open infrastructure that will achieve our joint vision of democratizing access for all.
- **We are encouraging all of you to join us!**
- To achieve our collective vision we can use all the help we can get ! Let's coordinate and collaborate.



Acknowledgements



- This work was partially supported by the NSF grants OAC-2112167, OAC-2030508, OAC-1841530, OAC-1836650, and MPS-1148698, CNS-1456638, CNS-1730158, ACI-150112, ACI-1541349, OAC-1450871 (NU), OAC-1659169 (SDSU)

